Gigabit Satellites in Distributed Supercomputing for Global Research

Larry A. Bergman

Jet Propulsion Laboratory

Abstract

This paper will examine the application of the NASA Advanced Communications Technology Satellite (A CTS) for implementing a distributed supercomputer global climate model, and in remote observation and control of the Keck telescope in Ha waii.

1: Introduction

1.1: Motivation

The ability to project high performance computing resources to remote corners of the globe will accelerate and in some cases, enable new science to be undertaken that otherwise would take years or even be impossible to perform. Three broad categories of satellite applications exist: (1) remote observation and control of distant science instruments, (2) distributed modelling and simulation, and (3) access to distributed data bases and archives.

1.2: Classes of applications

1.2.1: Remote **observation.** In many modern disciplines, such as geophysics, oceanography, climatology, petroleum research, seismology, astronomy, medicine, the ability to access and rapidly process large sets of previously recorded data can allow field measurements to be retaken as needed to improve accuracy. We expect that scientific discovery will greatly accelerate in many of the natural sciences as field scientists arc allowed to more rapidly compare and analyze field data with those previously logged in archives. Also, the ability to correlate or fuse between dissimilar data set types (e.g., thematic mapper, radar, elevation, seismic) can lead to new types of scientific discovery. Past and future missions that would benefit from this technology have been ozone investigations (in Antarctica), earth observing system (EOS–DIS), interplanetary missions, observation of solar eclipses, and volcanic and seismic surveys to name a few.

For a number of these important applications, high data rate instruments must be placed in locations that are inhospitable to human beings and/or arc difficult or impossible to connect with fixed high–speed networks. High rate (155--622 Mbit/s) satellites such as the ACTS permit for the first time these instruments to be controlled remotely from the convenience of one's office computer. In many cases, the instrument may be airborne, moving cm a truck, or transported on a ship at sea.

Onc of the key experiments discussed in this paper will explore the usc of ACTS to gather data from the world's largest optical and infrared telescope, the Keck Telescope located on Mauna Kea, Hawaii. Using the ACTS communications facilities to connect to the Keck telescope will make it possible to create a remote, full–function control room at Caltech, to enable individual researchers in California to carry out observations on the Keck Telescope just as if they were in Hawaii, and to perform real–time data logging from Keck Observatory into data archives. This project will be prototypical of many appl i cat ions that involve control and data acquisition from fast, remote instruments. The NASA Earth Observing System (EOS) is a premier example of such an application.

1.2.2: Distributed **modelling** and simulation. Another possible use of ACT'S is to setup temporary high speed networks between widely separated supcrcomputer centers to support specific distributed grand challenge applications, such as aerodynamic simulation, geophysics modelling, and global climate modelling (GCM). Over the past 15 years, high performance computing and communications (HPCC), and more recently meta-supercomputing, has becomc an essential tool in NASA's research and development program as a means to solve large scale scientific and engineering problems. In its usual form, this involves the decomposition of a problem solution into several subtasks, their assignment to individual computing resources and the interaction of the geographically distributed computing resources, data archives and visualization devices in order to achieve a reasonable speedup. Among its many advantages of satellite based meta-supercomputing [1] are:

- It quickly brings to bear the computational power of mot e than one machine and uses the most suitable machine for a specific subtask (a meta-supercomputer)
- It can access centrally located databases that may bc too costly or too impractical to replicate
- It can provide local visualization of results produced by ren iote applications programs.

1,2.3: Distributed data base access. A natural state of affairs in the scientific community is that centers of excellence tend to be dispersed globally—usually for political or mission charter reasons—that tend to attract specialists of a given discipline at specific research laboratories. Along with this, computational centers also tend to evolve that provide the custom libraries and computational scientists that solve the unique problems associated with that discipline. For example, oceanographic and radar libraries are archived at Jet Propulsion 1.aboratory (JPL), atmospheric libraries at Goddard Space Flight Center (GSFC) and Na-

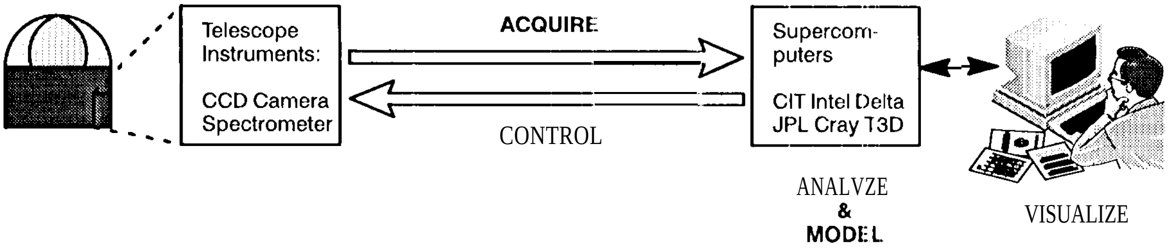

Fig. 1. Remote astronomy consists of four steps: (1) acquisition of the instrument data typically over a network, (2) analysis and modelling of the captured images, (3) visualization of the result, and finally, (4) adjusting the telescope parameters and retaking an exposure (control).

tional Center for Atmospheric Research (NCAR), seismic archives al several petroleum companies, and so cm. Since these data sets arc large (terabytes) and constantly being updated (in some cases), it is impractical to move them to one supercomputer site for onc application run. Hence, a distributed model offers the advantage of maintaining the locality of data as well harnessing additional CPU cycles.

2: Experiment objectives

Two major tasks are now proposed for the Experiment that will demonstrate the utility of ACTS of facilitating global science research. The first, called the *Keck Telescope Acquisition, Visualization and Control will* usc ACTS to control the Keck telescope from the Caltech campus in Pasadena, California. Astronomers will be able to usc the telescope more conveniently and perhaps even in an instructional setting. 10 some cases it may be desirable to usc the network connections to apply supercomputer resources to process or enhance images in real–time.

The second major task is the *Coupled Atmosphere–ocean Modelling for Global Climate.* In this case, two supercomputers, onc at GSFC and one at JPL/Caltech, will be temporarily connected by the high bandwidth ACTS channel to form a meta–supercomputer that will enhance the accuracy and accelerate the speed of global climate prediction, such as El Nine. This application demonstrates the feasibility and utility of using a gigabit satellite to setup temporary network connections between distance computer futures for special short lived projects or emergencies.

The kcy objectives of this ACTS experiment are to:

- Demonstrate distributed supercomputing (meta–supercomputer) over a high performance satellite link

- Demonstrate remote science data gathering and analysis with meta–supercomputer resources using multiple satellite hops

- Determine optimum satellite terminal/supercomputer host network protocol design to for maximum meta–supercomputer efficiency

3: Keck telescope

3.1: Overview

The Keck telescope acquisition, visualization and control experiment actually illustrates a form of telescience that might one day be attempted with NASA space missions. In telescience or remote observation, typically an investigator collects data (e.g., an image) with a remote instrument, examines the data with local computers, and then collects additional data with improved control parameters. Networks enable the use of instruments that are in locations unsuitable for humans, allow more scientists to use the facility, and permit more rapid analysis of data. Also, since ACTS would be linked into the CASA [2,3] gigabit fiber optic network in Pasadena, a variety of supercomputers could be used for compute-intensive tasks, such as enhancement of Keck telescope observational data in real–time (during the data acquisition phase).

3.2: Instrument data acquisition, visualization, and control

The Keck telescope will generate large volumes of astronomical data that must be transmitted to the astronomer as soon as possible so that he can optimize plans for subsequent observations and make the most effective use of limited telescope time. In this section, we describe the various phases of observing (Fig.1.), estimate their data rates, and describe tbe necessary response time for various modes of observing.

The initial set of general purpose scientific instruments for the Keck telescope, include two optical and two infra–red instruments. The low resolution imaging spectrometer (LRIS) is intended for observations of the faintest, most distant objects while the high resolution spectrometer (HIRES) offers a more detailed look at the spectra of somewhat brighter objects. Then there is a two micron imaging camera (N I RC) as well as a far infra–red ten micron camera (LWIRC). Additional more specialized instruments will be constructed for experimental use for particular problems such as interferometry.

3.3: Keck remote control software

All of the instrument and telescope control software has been written with the possibility of remote observing in mind. All commands are in the form of messages passed over an ethernet to either the telescope control computer or the instrument control computer. There is, of course, a subset of commands that for safety reasons can only be executed from the main telescope control console, as well as a larger subset of commands, used for engineering purposes only, which are restricted to those possessing the necessary passwords. Many X window telescope displays already exist and can be adapted to this application.

3.4: Keck data rate requirement

For illustrative purposes, we calculate the expected usage and data rates for the four major first light instruments. The table below (Table 1) lists the instruments, gives the percent of time that each is expected to be in use, the detector array size is given and it is assumed to be digitized 16 bits. The readout rate per pixel is also shown. The data rates are then, for the worst case, those corresponding to continuous readout, while data rates corresponding to more typical exposure times are also given.

The daily traffic volume from the family of telescope instruments is given in Table 2. These figures are estimated from the typical usage each evening, or for worse case, the maximum amount of data an instrument could generate per evening if run continuously. (For the Long Wavelength IR Camera, we assume the worst case is an interval of 1 second between exposures.) These numbers are total traffic in MBytes/night of observing. Each night is assumed to be 10 hours long.

The initial implementation of the charged coupled device (CCD) arrays uses 2048x2048 detectors and this will updated to 4096x4906 in the near future.

TABLE 1
Keck CCD camera readout characteristics

CCD CAMERA TYPE	% USE	CCD ARRAY SIZE	PIXEL READ TIME	INTERVAL BETWEEN FRAMES: Typical	INTERVAL BETWEEN FRAMES: Worse Case
LRIS	35	4096x 4096	4 usec	15 min	contin-uous
HIRES	20	4096x 4096	4 usec	15 min	contin-uous
NIRC	25	256x 256	20 usec	1 min	contin-uous
LWIRC	20	32x32	1 usec	1 min	contin-uous

Note – the 4 microsec/pixel on the optical arrays is actually a readout time of 16 microsec/pixel with 4 readout channels running simultaneously multiplexed together on the output.

continuous = very short exposures, with essentially continuous readout.

TABLE 2
CCD camera data volumes for Keck telescope

CCD CAMERA TYPE	TYPICAL NIGHT (MB)	WORST CASE (continuous operation)
LRIS	1,280	8,640
HIRES	1,280	8,640
NIRC	78	1,800
LWIRC	1.2	72

3.5: Keck remote calibration

The Keck instruments are complex and extensive calibration data is required to remove the instrumental signature in an optimum way. These calibration procedures are normally done in the late afternoon or early morning to save precious night time. The number of frames required is not excessively large, perhaps 100 or so, these data are not time critical and if the astronomer does not see them immediately after they are obtained, little is lost. But the astronomer does need to see them on a time scale of a half hour or so in order to ensure that the calibration procedure has worked correctly.

3.6: Keck remote observations

The deleterious physiological effects of the extremely high altitude of Mauna Kea make remote observing extremely attractive. Present plans call for observing from the sea–level headquarters in Kamuela over a T1 link (1 .55 Mbit/s) precisely to avoid this problem. A link to California via ACTS would reduce the time lost in travel, and offer convenient access to libraries, students and other members of the observing team. The data rates and volumes for normal observations with the Keck telescope have been calculated in Table 2 and 3, respectively. The astronomer would like to see the data as soon as possible to be sure that the observation was completed successfully. A small time delay of several seconds to a maximum of a minute is a tolerable price to pay for remote observing and is easily within the range of ACTS capabilities.

3.7: Keck instrument control

To fully implement remote observing, this ACTS link will also have to be used for telescope and instrument control. The traffic volume is negligible compared to that of the data. The only significant volume comes from the autoguider camera when one is finding a guide star or checking the identification of a field. If we assume frame rates updating at 10 frames/second, and a frame size of 256x256 digitized to 16 bits, the guider generates 0.13 Mbyte/frame, and a traffic of 1.3 Mbyte/see to be sustained over relatively short time interval not exceeding 10 minutes. Once a guide star is located, and everything is centered up, normal updates of the guider image will probably be at 5 to 10 second intervals and one could just transmit the guider error signals rather than the whole image.

Here, immediate response is critical and a delay in access to ACTS of more than a few seconds is fatal. However, this situation only holds during the few moments of frenzy at the beginning of the exposure. Thus, to summarize, during this critical phase, which will occur several to many times per night at unpredictable intervals, we need immediate access to ACTS.

3.8: Keck data logging

The present plans for archiving Keck observations call for storage of all raw observational data on Exabyte tapes. This is a compromise choice between factors of cost, ease of application, capacity of medium, medium storage life, etc. At present there is no plan for archiving reduced frames. The data rates and volumes are as given above. The actual archiving is probably best done in Hawaii to ensure the integrity of the data archive. Copies might be maintained at Caltech and made from transmissions of the satellite. It would be highly desirable to have a copy of the complete data archive on the continental U.S. (CON US) to enable full and easy utilization of this unique and valuable resource. The policies regarding access to the data archive and the time scale under which data is no longer the property of the original observer have not yet been established.

3.9: Keck ACTS demonstration

An interface to the telescope's image system and control system will be augmented from the existing DCS (Drive and Control System) system to allow the Keck telescope to be remotely controlled from the Caltech campus with minimal off-site operator assistance and permit data collection in near real--time at Caltech. An optional enhancement that will be explored if time permits is the use of large--scale computers at Caltech, JPL, and other CASA network sites to collect data automatically and make adjustments based on its content (e.g., cataloging celestial objects in sky searching).

3.10: Future

Eventually, JPL, GSFC, and Keck will be connected with low latency low bit error rate (BER) fiber optic cable. It is not our intent to demonstrate ACTS as a long lived communication service between these sites, but rather the opposite — that ACTS can be used to quickly assemble vast supercomputer resources and project them to remote regions around the globe in a relatively short time. This makes it a very powerful tool for investigators in the field to tap into satellite image data bases and supercomputers services. The application demonstrated here will be a model for a variety of field applications where fiber will normally not be available, such as providing remote science visualization and archival access for field expeditions (e.g., accessing LANDSAT, HIRIS, and SPOT data for archeological surveys, ozone hole investigation in Antarctica, seismic reflection profile trucks, and ocean floor sonar mapping ships). A department of defense (DoD) variation on this theme is the asynchronous transfer mode (ATM) "Global Grid" may be extended to remote points around the globe so that battlefield commanders can interactively explore high resolution satellite data in the theatre of operations during a regional conflict.

4: Distributed global climate modelling

4.1: Introduction

Coupled atmosphere-ocean general circulation models (GCMs) play a key role in the study of the climate system. These models explicitly solve the equations governing fluid motion on a rotating sphere, including parameterizations of physical processes at subgrid scales (e.g., cumulus convection, turbulent diffusion), and thus can be used to study nonlinear interactions and feedbacks between different components of atmosphere. and ocean circulations. Examples of outstanding problems studied with GCMS are El Nino–Southern Oscillation events and the impact on climate of increasing concentrations of greenhouse gases.

Currently, simulations using coupled atmosphere-ocean general circulation models (CGCMs) for climate studies (i.e., decades long) can only be performed at relatively low spatial resolutions because of their demands on computing resources. Simulations [4,5] with the University of California, Los Angeles (UCLA) CGCM using standard resolution require 30 CRAY Y- MP CPU hours per year at an execution rate of approximate y 130 Mflops. (The UCLA CGCM comprises the UCLA atmospheric GCM (AGCM) and the NOAA Geophysical Fluid Dynamics Laboratory (GFDL)/ Princeton University oceanic GCM (OGCM); standard resolution version of the AGCM is nine layers in the vertical with a grid spacing of four degrees latitude by five degrees longitude; standard resolution for the OGCM is 15 levels in the vertical and a grid spacing of one degree longitude by one degree latitude), Further, ensembles of simulations have to be conducted to assess the sensitivity of the climate system. Clear] y, long-term coupled simulations of one or more decades at higher spatial resolution require large speedups and high computational efficiency. Distributed computer environments connected by high–speed links have the potential to provide the necessary computer resources for climate research.

This project aims to explore the performance of a CGCM in a distributed computer environment consisting of a CRAY C98 at GSFC and a CRAY T3D at the NASA/Caltech/JPL, with the ACTS link. The application to be distributed has four main modules: 1) the UCLA AGCM, 2) the GFDL OGCM, the NASA/GSFC Aries AGCM, and 4) the NASA/GSFC Poseidon OGCMs. This coupling of GCMS will allow for scientific investigations on the coupled atmosphere/ocean system to be performed in collaboration between major NASA and University research centers. In the more specific framework of atmospheric and oceanic modelling, the distributed application will allow for research into the complex interactions and feedbacks of the coupled system in a broader context that is possible when only one AGCM or OGCM is used. Eventually, other modules can be added to the distributed application, such as those for chemistry and biological processes, as well as for data assimilation.

4.2: GCM/ACTS experiment definition

The task-decomposed coupled GCM has been run in a local heterogeneous environment consisting of a CRAY Y-MP at the Jet Propulsion Laboratory (JPL) and the Caltech Intel Delta connected by high performance parallel interface (HIPPI) network. Network communications between computers were handled by Express. Preliminary timings showed that the distributed model took 55 seconds per simulated day with a network communication overhead amounting to 3.6% of the total wall–clock time. The average communication bandwidth was 36.8 Megabits per second with a peak transfer rate of 63.2 Megabits per second. These transfer rates were limited by the speed at which Intel Delta communication nodes can import/export data. Using the ACTS link we will run the task–decomposed coupled GCM on the CRAY T3D at JPL and the CRAY C98 at GSFC.

4.3: Objectives

The overall research goal of this investigation is to explore methodology and performance issues for decomposing a coupled atmosphere-ocean GCM to run concurrently on heterogeneous computer architectures connected by a satellite link. Our research objectives include:

- Determining the communications bandwidth requirements for different decomposition strategies.
- Determining the effects of latency and communication cost for different decomposition strategies.
- Finding methods to mask latency with computation.
- Achieving a superlinear speedup of the coupled GCM code in a distributed, heterogeneous environment.

5: Network architecture

5.1: Network topology

The two science experiments will require a 3–node ACTS network (Fig. 2.) with HDR terminals installed at GSFC, JPL, and Hawaii (Tripler Medical Center, Oahu). The two links need not operate simultaneously. The GSFC/JPL link will be used for the GCM experiment while the JPL/Hawaii link will be used for the Keck telescope experiment.

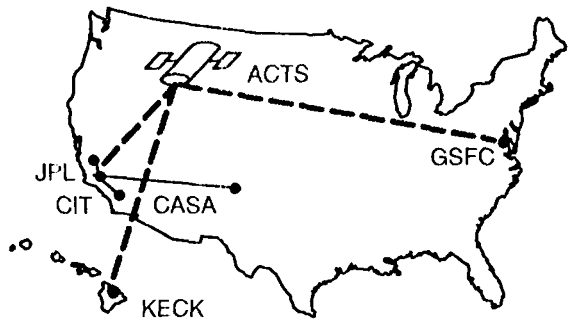

Fig. 2. The ACTS satellite experiment will consist of three principal nodes situated at JPL (Pasadena, CA), GSFC (Greenbelt, MD), and Hawaii (Honolulu or Mauna Kea). The JPL/GSFC nodes will be used for the GCM telescope experiment while the JPL/Hawaii nodes will be used for the Keck telescope experiment.

.2: ACTS network modelling

One key element of each application experiment consists of modelling the performance of the complete terrestrial and satellite network connection between the three HDR sites (JPL, GSFC, and Hawaii), and then using the model to predict the performance of the two applications (Keck and GCM). The approach is comprised of deriving the metrics for the applications (with the help of the application developers) and the network metrics, and then correlating the domains. Of particular importance are the effects of atmospheric burst errors, latency, latency variance, and the satellite channel multiplexing. Comparisons between the high gain and steerable antennas will also be made. Possible outcomes may be suggestions for implementation of particular feed forward error correction codes or the size of the packets for either best efficiency or control.

5.3: Supercomputer centers

The ACTS HDR will be interfaced to the supercomputer centers at GSFC and JPL/Caltech through a local ATM network. A backup approach will be to use a SONET/HIPPI interface adapter, Currently, the JPL/Caltech local computer network is based on HIPPI crossbar technology, operating at 800 Mbit/s. Nodes extend from the main computer center at JPL Woodbury (in Altadena) to both the JPL Oak Grove and Caltech campuses through fiber optic direct links. A portion of the CASA SONET OC-48 network is used between the JPL and Caltech. If necessary, CASA facilities at SDSC and LANL may also be used for these experiments (although these are not planned at the moment).

The Cray supercomputer interface to the ATM switch is a newly emerging product from Cray Research which will enable Cray supercomputers to connect natively into ATM–based networks at SONET OC-3, and later, OC-12. The initial product is referred to as a broadband gateway (BBG) and will connect to a standard Cray HIPPI channel. Eventually, a native ATM host interface will be introduced as well.

The ACTS 3.4-meter High Data Rate (HDR) ground stations will be placed near each respective institution's computer center. The interface consists of either four SONET OC-3 lines (155 Mbit/s) or one SONET OC-12 line (622 Mbit/s). SONET OC- 12 service is planned between JPL and GSFC. However, due to the fact that Hawaii is only serviced by the ACTS low-gain steerable antenna, only OC-3 service will be supported to the Hawaiian islands.

6: Conclusions

Gigabit pcr second satellite communications is expected to enable a number of new science and commercial applications over the next decade, and in some situations will offer key advantages over fiber optics—especially where rapid deployment and mobility are required by the user. Two emerging application areas that take advantage of this ncw capability are telescience and distributed computing. The higher latency and bit error rate of satellites is expected to pose some special problems for satellite applications, but are expected to be manageable with proper system architecture design. Working together with fiber optics, high rate satellites such as ACTS arc expected to make the global information highway a reality.

Acknowledgments

The author would like to gratefully acknowledge the invaluable support of key team members for this research including Burl Edelson and Neil Helm (GWU), Patrick Gary (GSFC), Paul Messina and Judith Cohen (Caltech), Carlos R. Mechoso (UC] .A), Jim Crisafulli and Bill Santos (PacSpace), and Ted Su and Benny Toomarian (JPL). The research described in this paper was performed by the Center for Space Microelectronics Technology, Jet Propulsion Laboratory, California Institute of Technology, and was sponsored by the National Aeronautics and Space Administration, Office of Space Access and Technology.

Bibliography

[1] Helm, N.R., H.J. Helgert, and B.I. Edelson: Supercomputer Networking Applications, AIAA international Communications Satellite Systems Conference, Washington, D. C., March, 1992.

[2 Blom, R. G., Bergman, L. A., Crippen, R. E., Frost, E. G., Hussey, K. J., Lyster, P. M., Okaya, D. A., and Stanfill, D. F., 'Interactive, Rcg,ional--Scalc Geological Data Exploration and Analysis Across a Gigabit Computing Network: A Part of the CASA Gigabit Network Testbed,' in the 9th Thematic Conference on Geologic Re.mote Sensing, Pasadena, CA, February 8–11, 1993.

[3] Lyster, P. M., Bergman, 1.., Li, P., Stanfill, D. Crippe, B., Blom, R., Pardo, C., and Okaya, D., 'CASA Gigabit Supercomputing Network: CALCRUST Three-dimensional Real–Time Multi–Dataset Rendering,' Supercomputing '92, Minneapolis, MN, November 17-20, 1993.

[4] Mechoso, C. R., C.–C. Ma, J. D. Farrara, J. A. Spahr and R. W. Moore, 1993: Parallelization and distribution of a coupled atmosphere–ocean general circulation model. Mon. Wca. Rev., 21, 2062-2076,

[5] Readhead, A.C.S., Nakajima, T.S., Pearson, T.J., Neugebauer, G, and Oke, J.B. 1988 Astronomical Journal Vol. 95, p. 1278.